CERN–RRB–2022-078
10 October 2022

Autumn 2022 Report of the Computing Resources Scrutiny Group

C Allton (UK), P Christakoglou (The Netherlands), A Connolly (USA), J Hernandez (Spain), J Kleist (Nordic countries), T Mkrtchyan (Germany), N Neyroud (France), P Sinervo (Canada, chair person), V Vagnoni (Italy), A Valassi (CERN, scientific secretary), J van Eldik (CERN)

1 Introduction

The Computing Resources Scrutiny Group (C-RSG) is charged with reviewing the requests of the four Large Hadron Collider experiments, ALICE, ATLAS, CMS and LHCb, for computing resources provided by the Worldwide LHC Computing Grid (WLCG), and making recommendations to the WLCG Resource Review Board (WLCG RRB) on allocations of these resources given the approved physics programs of the experiments. This report summarizes the Autumn 2022 C-RSG review of the computing resources being used for 2022, agreed-upon pledged resources for 2023 and expected utilization and preliminary projections for the 2024 WLCG year. Unless otherwise noted, the start of each reporting period is 1 April.

The information on current utilisation, the preliminary projections for 2024 requirements, subsequent questions from the C-RSG to each collaboration and face-to-face discussions with collaboration representatives for clarifications constitute the inputs examined to arrive at the C-RSG comments, findings and recommendations. This information was sufficient to allow the C-RSG to assess the preliminary 2024 resource estimates and to understand the contingencies and risks faced by the collaborations with respect to the need for and availability of computing resources.

In the following, the usual conventions for the WLCG computing model is used: T0 refers to the CERN computing resources, T1 refers to the Tier 1 facilities and T2 refers to the Tier 2 systems.

2 C-RSG Membership

The chair thanks the C-RSG members for their commitment and expert advice, and the experiment representatives for their collaboration with the C-RSG and this review process. Thanks are also due to the CERN management for its support and to the C-RSG scientific secretary, A Valassi (CERN) for the organization of this first face-to-face meeting in three years.

He has taken over this role from H Meinhard (CERN), who supported the C-RSG very well over the last decade. The members of the scrutiny group thank him for his professionalism, advice and the positive attitude he brought to this effort. We owe him a debt of graditude.

3 Interactions with the Experiments

The experiments were asked to submit their reports by August 30$^{\text{th}}$, 2022. The C-RSG thanks the experiments for the timely submission of their detailed documents [1–4]. The group acknowledges the computing representatives of the experiments for their availability, their constructive responses to the questions raised by the C-RSG and subsequent requests for further information. The dedicated face-to-face meetings with experiment representatives were particularly helpful and greatly appreciated by the C-RSG.

Specific sets of C-RSG referees were assigned to review the ALICE and LHCb requests. As usual, by agreement with ATLAS and CMS management, a single team of C-RSG referees scrutinized the ATLAS and CMS reports to ensure a consistent approach. The referees subsequently reported to the full C-RSG, which then developed the recommendations contained in this report.

In anticipation of the Spring 2023 scrutiny, the C-RSG asks the experiments to submit their documents for the April 2023 WLCG RRB by February 21st, 2023. The C-RSG requests that as part of their submission the experiments respond to the general and experiment-specific recommendations, as has now become practice. According to decisions taken in 2017 about the scrutiny process, the review reported here is the first opportunity to discuss resource estimates for 2024, allowing for further in-depth discussion with the experiments and LHC Experiments Committee (LHCC) during winter 2022/2023. The C-RSG plans to make recommendations to the WLCG RRB on the 2024 resource allocations after the Spring 2023 scrutiny.

4 Background for Preliminary 2024 Requests

The C-RSG met with the experiments at a time of significant uncertainty. The continued availability of Russian and Belarusian computing resources and network connectivity was uncertain. The LHC had recently had an unscheduled 4-week shutdown due to a cooling failure, raising questions of how much data would be collected in the 2022 calendar year. Finally, there was uncertainty in the longer term LHC schedule due to potential power limitations affecting the 2023 schedule. It was in that context that the collaborations discussed their preliminary estimates for computing resources for the 2024 year.

Each collaboration confirmed that the resources currently available in the 2022 year were sufficient to meet their physics goals. The approved resources for 2023 also appear to be sufficient to meet the expected needs of the collaborations, with perhaps some under-utilisation of tape storage resources due to the expected shorter calendar year 2023 LHC schedule. The computing plans for the collaborations anticipate the currently planned 2024 LHC schedule and concomitant data-taking schedule that then determined their preliminary resource estimates.

The C-RSG notes that the two experiments that are transitioning to new computing models, ALICE and LHCb, are making good progress albeit with the typical technical challenges that the commissioning of such large systems entail. The C-RSG congratulates these experiments in the progress that they have made and looks forward to updates on this progress in the Spring 2023 scrutiny round.

The C-RSG has become aware in recent years that the scaling used to justify a “constant-budget” model for annual increases of CPU and storage capacity may require adjustments for more current pricing and performance of replacement computing hardware. This development is being monitored by the WLCG and will help inform the constraints on these resources as the group considers computing resource planning.

In the following sections, the C-RSG summarizes the current computing plans and reports on the preliminary estimates for 2024, along with a number of findings and recommendations. The information provided for the ALICE and LHCb experiments is more detailed given their changing computing models and the effects on the expected computing requirements of both experiments.

5 Preliminary Resource Requests: ALICE

The C-RSG report is based on the usage and resource requests provided by the ALICE experiment [1], a written set of responses to C-RSG scrutiny questions, and a face-to-face meeting with the ALICE computing coordinators. The C-RSG thanks the ALICE team members for providing detailed responses to the C-RSG recommendations from the Spring 2022 scrutiny report and for their timely responses to questions from the C-RSG in this scrutiny.

The commissioning of the Online and Offline Computing (O2) system progressed this year with the stable beams being achieved by the end of May 2022. During this period ALICE was able to test the online calibrations and to utilize the Event Processing Nodes (EPNs) for synchronous processing. With the p-p data taking resuming at the end of September, ALICE expects to collect data at 500 kHz for 67 days resulting in 48 PB of Compressed Time Frame (CTF) output. For the commissioning runs the Time Projection Chamber (TPC) zero suppression was not running with its final dense format and the data rates are approximately 10 times higher than expected for the final format. If calibration of the TPC is not completed by the September run, ALICE has devised a selection strategy that should limit the total data size to about 10 PB which can be stored on the O2 disk buffer. With Pb-Pb data-taking scheduled for November 2022 (with 20 days of Pb-Pb data), completion of the calibration and dense format will be critical.

The current and estimated ALICE computing resources are summarized in Figure 1.

ALICE has been extremely successful in the acquisition of opportunistic resources in the first three months of 2022 with over 190 kHS06 from US supercomputing resources alone. These resources have been used to extend the physics channels that can be studied. The C-RSG congratulates ALICE on obtaining these resources. As noted in previous reports, once ALICE completes its commissioning and enters physics data-taking the C-RSG would appreciate the reporting of the usage of the O2 system (which represents a significant CPU resource).

The C-RSG notes that the pledged resources for T1 disk space continues to remain below the approved resource levels, but that the shortfall is not growing (standing at 90% of approved WLCG RRB resources compared to 85% in the previous scrutiny).

The C-RSG notes that the projected breakdown of disk and tape usage tracks ALICE's resource requests. The current utilization of these resources is, however, still at 73% disk and 54% tape capacity and will require careful monitoring.

Resource requirements for 2023 remain unchanged from the previous C-RSG endorsed numbers. ALICE expects to collect 128 days of p-p data at full magnetic field and 12 days of p-p data at low field, resulting in 52.2 PB of CTF output. The use of a more aggressive compression (Strategy B) will result in a reduction of these data to 0.6 PB. The low magnetic field data, for which all events will be kept in order to fully exploit the larger acceptance for soft observables, will lead to a further 6.4 PB of CTF output. At the end of 2023, 28 days of Pb-Pb data will collect 44 PB of CTF data (assuming Strategy B compression). The Strategy B compression (reducing CTFs from 56 PB to 44 PB) should be validated using the Pb-Pb 2022 data and will be reported on in the Autumn 2023 scrutiny (prior to the 2023 Pb-Pb data-taking).

Reconstruction of the 2023 Pb-Pb data will be undertaken in 2024 using two asynchronous passes spread equally across the O2, T0 and T1 resources and lasting for a period of 10 months. The data and CPU requirements for simulations to support the Pb-Pb runs remain based on Run-2 parameters. For data taking in 2024, ALICE anticipates 140 days of p-p data at 500 kHz and 28 days of p-Pb. In addition, two short O-O and p-O runs are expected to be undertaken. The Strategy B compression is expected to be in place for these runs, which will result in approximately 49 PB of CTF data. As before, compute times for the asynchronous reconstruction and the Monte Carlo (MC) simulations are based on predictions from Run-2 benchmarks. With the data challenges completed over the summer

ALICE		2022		2023			2024	
		C-RSG recomm.	Pledged	Request	2023 req. / 2022 C-RSG	C-RSG recomm.	Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	471	471	541	115%	541	622	115%
	Tier-1	498	448	572	115%	572	655	115%
	Tier-2	515	517	592	115%	592	683	115%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**1484**	**1436**	**1705**	**115%**	**1705**	**1960**	**115%**
	Others							
Disk	Tier-0	50.0	50.0	58.5	117%	58.5	67.5	115%
	Tier-1	55.0	49.7	63.5	115%	63.5	71.5	113%
	Tier-2	49.0	55.2	57.5	117%	57.5	66.5	116%
	Total	**154.0**	**154.9**	**179.5**	**117%**	**179.5**	**205.5**	**114%**
Tape	Tier-0	95.0	95.0	131.0	138%	131.0	167.0	127%
	Tier-1	63.0	71.8	82.0	130%	82.0	102.0	124%
	Total	**158.0**	**166.8**	**213.0**	**135%**	**213.0**	**269.0**	**126%**

Figure 1 Alice resource requests and C-RSG recommendations for 2022 and 2023, and resource estimates for 2024.

of 2022 and the data taking expected before the next scrutiny, the C-RSG would appreciate a report on the comparison of the actual times required to undertake the p-p and Pb-Pb simulations compared to the times originally predicted from the Run-2 benchmarks. This will allow a better understanding of the resource requirements for the p-Pb simulations.

For 2024, ALICE estimates result in a growth of +15% for CPU (across all tiers) and +14% for disk (divided 15%, 13% and 16% across T0, T1 and T2, respectively). To archive the CTF files requires a +26% growth in tape (divided 27% and 24% across T0 and T1, respectively).

If the Russian and JINR resources are not available (or compensated for) by the funding agencies the resource requests grow to +24% for disk (divided 15%, 32% and 26% across T0, T1 and T2, respectively) and 25% for CPU (divided 15%, 34% and 26% across T0, T1 and T2, respectively) Without the Russian or JINR tape resources the tape request will grow by 34% (divided 27% and 45% across T0 and T1, respectively).

The 2024 disk requests are driven by the processing of the 2023 Pb-Pb data (see above) and approximately half of the first reconstruction of the 2024 p-Pb data with 15 PB for Pb-Pb processing and 14 PB for p-Pb. Disk resources are modest for p-p processing (2023 and 2024) at 2 PB and for O-O and p-O processing only 1 PB. Removal of Run-2 data (in Run-2 format) should reduce the disk requests by approximately 6 PB.

In 2024, 30% of the CPU capacity will be devoted to analysis workflows, while 20% of the CPU capacity will be dedicated to processing the 2023 Pb-Pb data and reconstruction of the 2024 p-p, O-O, p-O and p-Pb data. By the end of 2024, 50% of the CPU capacity will be for p-Pb MC production.

Conclusions and Recommendations

The C-RSG compliments the ALICE experiment on its work on commissioning the O2 system.

ALICE-1 The C-RSG suggests that for future Autumn scrutinies ALICE reduce the length of the report by focusing on the requests for future computing resources and a justification for these resources. The C-RSG recommends reducing the detail provided in Section 2 and excluding Section 3.

ALICE-2 To evaluate the robustness of the predictions for 2024 resource requests, the C-RSG asks that for the next scrutiny ALICE provide a comparison of the predicted performance of their simulations (in terms of required CPU and disk) that were derived from Run-2 benchmarks with the actual performance that is being achieved using the O2 framework today.

6 Preliminary Resource Requests: ATLAS

The ATLAS report is based on the information provided by ATLAS [2], written responses to questions by the C-RSG, and a hybrid face-to-face meeting with the ATLAS computing coordinators.

The current and estimated ATLAS computing resources are summarized in Figure 2. The new Run-3 data analysis format, DAOD_PHYS, has proven successful and well below its target size of 50 kB/event for MC and 30 kB/event for data. For Run-4 a new data format, DAOD_PHYSLITE, which is 2-5 times smaller than DAOD_PHYS, is under adoption. The first analyses using these formats should be finalized in the middle of 2023. A large part of the 2024 production is expected to make use of DAOD_PHYS and DAOD_PHYSLITE.

The new fast simulation, AtlFast-3, is being used for Run-2 needs. Early next year in-depth studies are expected to reveal whether the physics performance goals will be reached while keeping the required CPU needs at the level of the previous version, AtlFast-2. This will also clarify whether the goal of producing 60% of the simulated Run-3 events using the fast simulation in 2024 is achievable.

The distributed, decentralized ATLAS computing model does not require that specific sites be available. Thus a temporary unavailability of any CPU resource at a single site could be absorbed, in the worst case, by delaying some of the analyses. The loss of disk and archival storage (tape) is extremely difficult to mitigate directly. Instead, ATLAS is working with a number of new projects to provide archival storage outside T1 sites. Should this R&D prove successful, the loss or restriction of tape capacity at some T1 sites could be mitigated at least temporarily.

The ATLAS Collaboration makes the following assumptions on the running scenario for 2024: one full year of operation with an integrated luminosity of $L = 100\ \text{fb}^{-1}$, pile-up of 50 collisions per bunch crossing and a total running time of $t_{run} = 6 \times 10^6$ s for p-p and 1.2×10^6 s for heavy ion runs. With these assumptions, ATLAS expects to record 10.2×10^9 events in the main trigger stream and an additional 9.6×10^9 events in a delayed stream.

ATLAS computing requests for 2024 are driven by the following planned computing activities:

- processing and reprocessing of Run-3 data taken in 2024 using the newest release;
- simulation of MC samples for 2024 (campaign starting in 2023);
- starting the simulation of MC samples for 2025;
- production of derived formats for data and MC (DAOD_PHYS and its skims as well as DAODs made directly from AOD);

ATLAS		2022		2023			2024	
		C-RSG recomm.	Pledged	Request	2023 req. / 2022 C-RSG	C-RSG recomm.	Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	550	550	740	135%	740	850	115%
	Tier-1	1300	1300	1430	110%	1430	1501	105%
	Tier-2	1588	1588	1747	110%	1747	1834	105%
	HLT	n/a	n/a	n/a	n/a	n/a	0	n/a
	Total	**3438**	**3438**	**3917**	**114%**	**3917**	**4185**	**107%**
	Others							
Disk	Tier-0	32.0	32.0	40.0	125%	40.0	46.0	115%
	Tier-1	116.0	116.0	136.0	117%	136.0	162.0	119%
	Tier-2	142.0	142.0	168.0	118%	168.0	198.0	118%
	Total	**290.0**	**290.0**	**344.0**	**119%**	**344.0**	**406.0**	**118%**
Tape	Tier-0	120.0	120.0	174.0	145%	174.0	205.0	118%
	Tier-1	272.0	272.0	353.0	130%	353.0	448.0	127%
	Total	**392.0**	**392.0**	**527.0**	**134%**	**527.0**	**653.0**	**124%**

Figure 2 ATLAS resource requests and C-RSG recommendations for 2022 and 2023, and resource estimates for 2024.

- user analysis on Run-2 and Run-3 data and MC samples;
- continuation of physics studies for the HL-LHC phase; and
- processing heavy-ion data and MC.

The experiment is asking for increases of 6% for CPU, 18% for disk and 24% for tape, justified by the expected data collection in 2023 and 2024.

Conclusions and Recommendations

ATLAS-1 The C-RSG recommends that ATLAS continue increasing the adoption of the more compact DAOD_PHYS and DAOD_PHYSLITE data formats for its physics data analyses.

7 Preliminary Resource Requests: CMS

The CMS report is based on the information provided by CMS [3], written responses to questions by C-RSG, and a hybrid face-to-face meeting with the CMS computing coordinators. The current and estimated CMS computing resources are summarized in Figure 3.

The updated CMS detector is performing in 2022 as expected. Analysis has already begun on the acquired data, starting from samples in NanoAOD format, demonstrating the reliability of the format and the related tools that CMS builds and distributes.

CMS		2022 C-RSG recomm.	2022 Pledged	2023 Request	2023 req. / 2022 C-RSG	2023 C-RSG recomm.	2024 Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	540	540	720	133%	720	750	104%
	Tier-1	730	730	800	110%	800	860	108%
	Tier-2	1200	1200	1350	113%	1350	1500	111%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**2470**	**2470**	**2870**	**116%**	**2870**	**3110**	**108%**
	Others							
Disk	Tier-0	35.0	35.0	45.0	129%	45.0	52.0	116%
	Tier-1	83.0	83.0	98.0	118%	98.0	108.0	110%
	Tier-2	98.0	98.0	117.0	119%	117.0	130.0	111%
	Total	**216.0**	**216.0**	**260.0**	**120%**	**260.0**	**290.0**	**112%**
Tape	Tier-0	155.0	155.0	228.0	147%	228.0	293.0	129%
	Tier-1	260.0	260.0	316.0	122%	316.0	370.0	117%
	Total	**415.0**	**415.0**	**544.0**	**131%**	**544.0**	**663.0**	**122%**

Figure 3 CMS resource requests and C-RSG recommendations for 2022 and 2023, and estimates of computing resources for 2024.

As a result of a coordinated effort of the physics and computing teams in CMS, the foreseen tape cleanup campaign took place in June and July 2022. About 20% of the total CMS data on tape have been deleted, mainly comprised of the non-legacy processing of Run-2 data and MC sample. This was the largest deletion campaign ever performed by the experiment and the first one since the adoption of the RUCIO data management system.

Since the Run-3 trigger farm has been installed and is operational, some of the Run-2 HLT resources are available as a permanent cloud for offline processing. By 2024 these machines will be out of warranty but still quite capable, allowing CMS to reduce CPU requests across all Tiers.

In 2022 CMS experiment has switch from gzip to Lempel-Ziv-Markov (LZMA) algorithm for the compression of raw data. The expected additional ~4% CPU overhead at T0 is aimed to be compensated by ~10% reduction of the RAW event size, which reduces tape storage requirements at T0 as well as at the T1 facilities. Such optimization might have a positive effect on the network utilization as well.

The flexibility of the CMS computing model mitigates the loss of CPU resources at a single site, including the loss of CPUs at an entire T1 facility, by redistributing the workloads or, in the worst case, by delaying some of the reconstruction or calibration activities. On the other hand, the loss of T1 storage, especially tape, is harder to compensate. CMS has noted that by having only six T1 sites, the loss of one of those for a period longer than a week would severely impact the ability to store and reprocess the experimental data.

CMS expects that 2024 will be a full data-taking year with the machine performance being very similar to 2023: integrated luminosity of $L = 100\ \text{fb}^{-1}$, pile-up of 50 collisions per bunch crossing and a total running time of $t_{run} = 6 \times 10^6$ s for p-p and 1.2×10^6 s for heavy ion runs. The choice of p-Pb instead of Pb-Pb for the heavy-ion run has a marginal impact on the CMS computing needs.

CMS foresees the following processing activities in 2024:

- Run-3 data-taking;
- Run-3 MC production;
- HL-LHC MC production; and
- Run-3 user data analysis.

The experiment is asking for increases of 8% for CPU, 12% for disk and 22% for tape. These are well-justified by the expected physics data-taking.

Conclusions and Recommendations

CMS-1 The C-RSG applauds CMS for their continuous efforts in making their software and computing environment more efficient in order to minimise their resource needs.

CMS-2 The C-RSG recommends to the CMS collaboration that it investigate the economic trade-off between raw data compression CPU overhead versus the reduction of tape storage requirements at T0. This information would be helpful for understanding the overall benefits of such a data compression strategies.

8 Preliminary Resource Requests: LHCb

The assessment of the preliminary LHCb computing resource requests for 2024 is based on the document provided by the experiment and discussions between the C-RSG and LHCb computing management. The initial 2024 resource request document was expanded [4] as a result of clarifications requested by the C-RSG.

Figure 4 shows the LHCb resource projections for 2024, as well as the C-RSG recommendations for 2023 and the pledged resources by funding agencies in 2022. At the T1 level, a deficit of pledged resources for 2022 of 17%, 10% and 16% for CPU, disk and tape, respectively, can be observed relative to the C-RSG recommended capacities. The corresponding deficit at the T2 level reaches 4% for CPU and 32% for disk. Due to the delayed start-up of the LHC Run-3 and the ongoing commissioning of the upgraded detectors, this shortage is expected to have no effect on the experiment provided the requested resources for 2023 are delivered on time.

However, a similar shortfall of pledged resources is observed for 2023 in the CRIC accounting portal [5]. This deficit could compromise the computing activities of the experiment planned for 2023. The installation of the LHCb upgraded detector is expected to be completed during the 2022/23 winter shutdown and the detector commissioning completed by the start of physics data-taking in 2023. LHCb will reassess the situation in the next scrutiny round in Spring 2023, evaluating any potential impact on the resource requests for 2024.

In 2024, LHCb projects large increases, close to 60%, in CPU, disk and tape capacities. This growth, in accordance with the LHCb computing model, is requested to store and process the experimental data to be recorded in 2024 and to fulfil the growing data simulation needs of the experiment. CPU requirements are dominated by MC simulation for Run-3 data, consuming close to 75% of the total LHCb CPU budget, out of which 4% will be used for heavy ion MC events. LHCb is not expecting to produce simulations for Run-1 and Run-2 conditions in 2024. The CPU needs for Run-3 simulations

LHCb		2022 C-RSG recomm.	2022 Pledged	2023 Request	2023 req. / 2022 C-RSG	2023 C-RSG recomm.	2024 Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	189	189	215	114%	215	344	160%
	Tier-1	622	515	707	114%	707	1128	160%
	Tier-2	345	333	391	113%	391	630	161%
	HLT	50	50	50	n/a	50	50	n/a
	Total	**1206**	**1087**	**1363**	**113%**	**1363**	**2152**	**158%**
	Others							
Disk	Tier-0	26.6	26.5	30.3	114%	30.3	46.8	154%
	Tier-1	52.9	47.8	60.5	114%	60.5	93.6	155%
	Tier-2	10.2	6.9	11.6	114%	11.6	18.1	156%
	Total	**89.7**	**81.2**	**102.4**	**114%**	**102.4**	**158.5**	**155%**
Tape	Tier-0	81.0	81.0	91.0	112%	91.0	188.5	207%
	Tier-1	139.0	116.0	157.0	113%	157.0	199.7	127%
	Total	**220.0**	**197.0**	**248.0**	**113%**	**248.0**	**388.2**	**157%**

Figure 4 LHCb resource requests and C-RSG recommendations for 2022 and 2023, and estimates of computing resources for 2024.

in 2024 have been calculated assuming no speed-up with respect to the 2023 simulation performance. Simulation parameters for 2023 (same event processing time as for Run-2 conditions in spite of a higher event complexity in Run-3) are considered optimistic by LHCb, and thus a conservative estimation for 2024 has been used. In fact, the Run-3 simulation is currently a factor 1.7 slower than the Run-2 simulation. The removal of already identified hotspots in the Run-3 simulation should bring its performance closer to that of the Run-2 implementation. LHCb will report on these expected improvements in subsequent scrutiny rounds.

The simulation of p-p collision data will ramp up to the nominal level for 2023 (4.8×10^9 events per fb^{-1} of collision data collected), while for 2024 this number will be at 50% of the nominal level. In addition, the LHCb Collaboration is expecting a significant increase in the pile-up rate by roughly a factor of six in Run-3. The C-RSG notes that the simulation of pile-up collisions in the LHCb full simulation is currently done by generating every underlying minimum-bias collision together with every signal collision. Given the significant increase in CPU requirements, the C-RSG encourages LHCb to investigate the possibility of combining signal with minimum-bias underlying events, the latter extracted from a pre-existing pile-up library. This approach, followed by CMS and ATLAS, results in large savings of CPU time due to the fact that the same minimum-bias events can be reused multiple times.

The fractions of events produced with full, fast and parametric simulations in 2024 are assumed to be the same as in 2023 (36%, 64% and 0%, respectively). Fast simulation is about 10 times faster than full simulation. Parametric simulations are expected to be used once a detailed understanding of the detectors and their calibration is established. The CPU work for user analysis is expected to scale with the CPU work for data stripping, in a similar way as for Run-2, albeit with a reduction factor of

50%. This stems from the transition of the majority of user analysis to a centrally managed analysis production system in conjunction with an analysis framework that is reorganised with a focus on CPU performance.

The projected increase in data storage is driven by the integrated data volume expected to be recorded in 2024. An additional 140 PB of tape storage will be needed to store two replicas of the raw data, and close to 55 PB of disk storage will be required to keep on disk two replicas of the latest data reconstruction pass for physics analysis. No significant reduction of these requirements is expected, neither by data format reductions nor reductions in the number of replicas. The average event size of the analysis data formats in Run-3 after content slimming is assumed to be around 35 kB. The C-RSG encourages LHCb to explore the possibility to adopt even more reduced data formats for analysis as a way to reduce the disk storage requirements.

The C-RSG notes the advances the LHCb Collaboration has made to integrate and exploit HPC opportunistic resources at various centres, e.g. Spain, Italy and the United States. Nevertheless, the projected level of utilisation of HPC resources in 2024 has not increased. The C-RSG encourages LHCb to continue identifying increased allocation opportunities in HPC centres.

The C-RSG also notes the significant effort the collaboration has made to address its apparent shortage of personnel available for computing activities. An effort equivalent to 2.2 FTEs has been obtained, mostly provided by CERN. While it acknowledges the significant existing commitment of personnel, the C-RSG encourages the other LHCb institutions to increase their contributions to computing given the continued shortfall in personnel.

Conclusions and Recommendations

LHCb-1 In view of the completion and commissioning of the upgraded LHCb detector, the C-RSG expects the experiment to reassess its 2024 resource estimates in the next scrutiny round.

LHCb-2 The experiment reported that the full simulation of Run-3 events is currently significantly slower than the expected performance. The C-RSG asks LHCb to report in the upcoming scrutiny round on the actual performance of the Run-3 simulation compared to the assumed performance for the 2023 and 2024 CPU requests.

LHCb-3 Simulation of pile-up collisions in the full simulation is currently done by generating every underlying minimum-bias collision together with every signal collision. Considering that expected CPU requirements are dominated by full MC simulation for Run-3 data, and that the level of pile-up increases by a factor of six from Run-2 to Run-3, the C-RSG recommends LHCb explore alternative pile-up modelling strategies to reduce CPU requirements.

LHCb-4 In view of the projected increase in data storage, the C-RSG encourages LHCb to explore more reduced data formats for analysis as an avenue to significantly reduce the disk storage footprint.

LHCb-5 Despite the progress the LHCb Collaboration has made in integrating and exploiting HPC opportunistic resources at various centres, the CPU time contributed by HPCs in 2024 is not planned to increase relative to the current utilization level. The C-RSG encourages LHCb to continue looking for increased allocations in HPC centres.

LHCb-6 The C-RSG congratulates the LHCb Collaboration for the efforts made to address the shortage of personnel available for computing activities. The C-RSG reiterates the importance of having additional LHCb institutions contributing increased efforts to computing.

9 Comments and General Recommendations

The C-RSG makes the following observations and recommendations regarding the status of computing resources for the year 2023 and the development of the resource requests for the year 2024.

ALL-1 The C-RSG notes that there may be reductions in pledged resources due to power availablity. It recommends that the WLCG collaboration identify the risks arising from this development and, where appropriate, develop contingency plans.

ALL-2 Several experiments will be placing increased demands on T1 tape resources. The C-RSG recommends that an assessment be made of the effect on tape storage capability of another loss of a T1 facility.

ALL-3 Given the shorter 2023 data-taking period, the C-RSG recommends all experiments review the effect this has on the estimated 2024 resources and report on any changes that arise from the LHC schedule adjustments in the Spring 2023 scrutiny.

References

[1] ALICE Collaboration. ALICE Computing Resources Usage January to July 2022, Resource Requirements for 2023 and Estimates for 2024, submitted 30-Aug-2022.

[2] ATLAS Collaboration. ATLAS Computing Status and Plans – Report to the C-RSG September 2022, submitted 30-Aug-2022.

[3] CMS Collaboration. Fall 2022 Scrutiny – CMS Preliminary Resource Request for 2024 - CMS Offline Software and Computing, submitted 30-Aug-2022.

[4] LHCb Collaboration. LHCb Computing Resources: preliminary 2024 requests. *LHCb-PUB-2022-016*, submitted 30-Aug-2022.

[5] WLCG. CRIC list of pledges.